Proceedings of the 2008 Winter Simulation Conference
S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler eds.

MMOHILS: A SIMPLER APPROACH TO VALID AGENTS IN HUMAN SIMULATION STUDIES

Seth N. Hetu
Gary Tan

Dep. of Computer Science
Computing 1, Law Link
National University of Singapore
Singapore, 117590

ABSTRACT

A novel technique for accurately and inexpensively simulating large numbers of people is introduced: Massively Multiplayer Online Human In the Loop Simulation (MMOHILS). This technique is applicable to certain simulations which would normally use AI-based agents to model human behavior, and its validation techniques are substantially simpler. A prototype for two types of MMOHILS (experimental and unannounced) is laid forth in this paper, with examples given from a current prototype in development for use in egress analysis.

1 INTRODUCTION

Classic simulation becomes expensive and challenging when humans are included in the model. Simulated computer agents can be built for well-studied domains such as pedestrian movement, but extending these agents to act realistically in more complex environments is often very difficult. Moreover, observation of human experiments or of recorded footage is often infeasible —for a variety of practical and ethical reasons— in extremes such as crisis situations. This paper provides a prototypical model for simulating real humans in a way that is both valid and inexpensive. In order to do this, a brief survey of existing approaches to software agents and live experiments is given. After stressing some of the shortcomings of each approach, a new method is put forth. Our prototype uses real human agents instead of software agents, and validates the remainder of the model with respect to the agents' interaction with their virtual world. The resulting system represents human agents virtually in a way that is inexpensive, valid, and flexible enough to allow substantial amounts of novel research to be carried out.

The remainder of this paper is organized as follows. In Section 2, we discuss previous endeavors and cite specific surmountable weaknesses in each approach. Section 3 discusses our prototype, with special emphasis on validation techniques. Finally, Sections 4 and 5 conclude the paper, after a brief discussion of some additional benefits our system enjoys.

2 PREVIOUS WORK

Solutions involving a human component have usually proceeded in one of two ways. First, in certain cases a controlled experiment with real subjects can be carried out. If this is infeasible, a second approach is to combine observations of real agents with sophisticated decision-making models, and to run a computer simulation.

On a slightly tangential note, some research exists which uses gaming and simulation for training, (Jain 2006) being a good example. Our approach is wholly different, using gaming to benefit the simulation and not the other way around. However, some of the techniques are the same.

2.1 Controlled Experiments

An abundance of experimental data exists in the domain of pedestrian movement. Daamen (2003) is a good example of a well-structured experiment. Most of these experiments are kept simple; Daamen's work deals with cross-flow, lane formation, and bottlenecks in small (less than $10m^2$) fields. These are also usually ideal studies: cameras are positioned directly above the field, and participants are instructed to follow carefully-worded directions such as "walk slowly and continuously from Point A to Point B" or "walk forward slowly and pause periodically."

Daamen's work concluded that, under conditions of *pedestrian* flow, participants filled the available width, even when pedestrian density was low. Lanes formed wherever possible, and lasted for a short duration. Cross-flows resolved when single pedestrians gave way to groups of two or more walkers. Pedestrian flow fanned out slightly with the introduction of a bottleneck.

Work such as Daamen's provides valuable insight into pedestrian flow, but experimentation in general has limited applicability. Un-announced observations such as those carried out by (Fruin 1987) and the Legion team (Berrou 2000) attempt to observe situations as they naturally occur. Observation usually suffers from less-than-ideal conditions —cameras, for example, usually cannot be mounted directly overhead and their input must be later corrected for distortion. Moreover, although observing pedestrians is possible, observing other human behaviors, such as egress, can be more difficult.

2.2 Simulation with AI Agents

When observation and human experimentation are infeasible, an appeal is often made to simulation. Chaturvedi's work provides a good case study in this regard. A night club fire was simulated, with the intent of determining how the geometry of the building affected the number of deaths. (Chaturvedi et al 2006) There were no recordings of the incident, and no reasonable experiment could be designed to reproduce the levels of stress and fear that factored heavily in this event's outcome. Instead, a simulation was constructed. The simulation had two components: a fire model and an agent model. Each agent was given an imperfect amount of knowledge about the building's structure; agents chose the most familiar and safest exit as a target exit and moved towards it according to a force-based model. The two components interacted with each other: opening windows affected the flow of the fire, and fire spread caused agents to change their target exit.

Okazaki's work uses magnetic forces to simulate goal-seeking and collision avoidance; however, in that simulation goals were statically assigned. (Okazaki 1993) A common criticism of this approach is that it "equips the individual occupant with certain specialized features pertaining to movement, but not with calculations relating to social capacities." (Santos 2004) For example, a force model will optimize exit utilization ("queuing") whereas a rule-based system such as Pan's will more often focus on social aspects of human decision-making (e.g., competitiveness, leadership) which lead to exit under-utilization ("herding"). (Pan et al 2005)

Unfortunately, there is no guarantee of validity as simulations with AI agents scale. Even Pan's simple observation that exit utilization is not ideal brings into question the validity of a great many otherwise-well-founded simulation studies. Yet validation of rule-based systems may prove too challenging and costly for some circumstances. What is needed is an alternative to complex AI-based agents that is both easy to validate and inexpensive to employ. The purpose of such a method is to enable simulation of human behavior under circumstances previously deemed too difficult, dangerous, or uncertain for traditional methods.

3 MASSIVELY MULTIPLAYER ONLINE HUMAN-IN-THE-LOOP SIMULATION

There is a phrase: "making mountains out of molehills", which means turning a simple problem into something needlessly difficult and complicated. We feel this is an accurate description of some of the problems faced when modeling AI agents. Certainly, there are many cases where complex agents are necessary; however, we argue that there exists a much simpler alternative.

Taken broadly, the main difficulty in modeling humans is in capturing the "human factors" that influence decision making. A user may choose a sub-optimal exit because she is not familiar with the faster route, or perhaps because she would rather follow someone who seems to know where he is going. Knowledge of these kinds of factors is often impossible to obtain *a priori*, and thus greatly complicates modeling and validation.

Human-in-the-loop simulation is a technology often used for, say, training a human aircraft pilot. The primary goal is to benefit the pilot by allowing her to train on all possible operating conditions. However, the simulation also benefits from the pilot; her reactions can be recorded, analyzed, and incorporated into the model of, say, an AI-controlled pilot.

Given this, we propose combining the strengths of Massively-Multiplayer Online experiences with Human-In-the-Loop Simulation to form a Massively Multiplayer Online Human-In-the-Loop Simulation (MMOHILS). A large number of users will be tracked as they progress through a simulation, and will provide a valid "human factor" which might otherwise be missing from AI-based methods. The remainder of the model can be validated using traditional methods. Moreover, any "human factors" can be validated using psychological evaluations such as surveys. (If, for example, we wish to validate the users' actions under stress, we might ask them to rate how stressed they felt after the experiment.) We propose two types of MMOHILS: experimental and unobserved. Both of these will be detailed, and then validation techniques, which are common to both methods, are discussed. We choose to pronounce MMOHILS as "molehills", to keep with our earlier analogy.

3.1 Experimental MMOHILS

Experimental MMOHILS is applicable to large groups of users all of which are knowingly participating, with good intention, in an experiment. Unlike classical simulation, one must consider the agents' individual knowledge of the virtual world. To achieve a proper confidence estimate, there should be several sessions with identical setups, possibly breaking larger groups into smaller sessions if the novelty of the environment is an important factor.

Discrete, time-stepped simulation is used to model user input, as it co-operates nicely with a technique used in online gaming known as the "lock-step" protocol. Using this protocol, the central simulation waits until all client programs have submitted input, and then processes this input. Next, the simulation issues to each user (client) a message containing the results of that input. Each client then proceeds with the next user input. A "maximum" lockstep value should be determined based on known estimates of network lag; any clients which do not update within one lockstep should be removed from the simulation.

Experimental MMOHILS is suitable for local area networks (LAN), which feature round-trip network delays of 1 to 2 ms. (Tierney 1994) Moreover, if all computers connected to the LAN can be monitored (e.g., they are all in the same lab) then it is unlikely that users will be able to cheat through modification of the game client. This is ideal, because it means that the client programs can perform the bulk of the computations at each time-step, and the central simulation (server) will only need to receive updates and issue lock-step messages.

Allowing client programs to process the bulk of the data, and assuming a 2ms round-trip delay, a great deal of complex virtual worlds may be developed. The perspective could be fully 3-dimensional, but this is not a requirement. For example, simulating evacuation of a building may demand only a top-down perspective.

3.2 Unannounced MMOHILS

Using MMO games as a research platform has been suggested before, (Tay 2005) but there are questions of how to retain users. Moreover, the goodwill of the users cannot be assumed, as is the case with experimental MMOHILS. Therefore, data gathering in a massively-multiplayer environment is best accomplished with techniques from the domain of unannounced observation. Players may suspect that they are part of an experiment, but they have no way of knowing if this is true nor, most importantly, can they ascertain what they are being evaluated on. They are, in a manner of speaking, "blind" to the output parameters.

Unannounced MMOHILS contains a number of hurdles that its experimental counterpart does not. Motivation is the first issue: how do we convince online game players to take part in our virtual world? Essentially, a MMOHILS functions on hosting "events" with "rewards". An "event" is an MMO term referring to a specific, announced time in the game world at which all logged-on users are transported to a separate map for a specific competition. Events feature "rewards" which can be used/worn in the game world at any time, and which typically are not found in the non-event sections of the game world.

Of course, given sufficient funding, cash prizes might be awarded; however, this approach is fraught with difficulty. For one thing, it is best to reward all users, so that they will intend to finish the simulation once they begin. Small monetary incentives may not provide that incentive. Instead, we propose rewarding all users who complete the event with in-game items which can only be found during events, and hosting events at the same time every day. An undisclosed number of these events will be placebos, and the rest will be runs of our simulation, using participants as live input. At all non-event times, the game world is running normally, and no simulation can be performed (i.e., there is no need to validate the "fun" parts of the game, which exist only to entice users to join the online community.)

An example might help to clarify this. Suppose that, in non-event mode, players can explore the virtual world, fight each other, or cast magical spells. (During event mode, battles and spells are disabled.) Fighting reduces an opponent's "vigor"; a battle is won when one player's vigor reaches zero. Spells reduce the caster's "vim", and perform a variety of effects. Both vim and vigor replenish gradually over time, and they can be restored to 100% immediately by consuming "vim potions" and "vigor potions" respectively. This setup induces a huge demand for vim and vigor potions; they allow players to avoid tiresomely waiting for vim and vigor to replenish naturally and "get back in the game" immediately. An example of an event in this case might start by transporting all players to a new location, and then informing them that they will be rewarded with 10 vim and 10 vigor potions if they reach a specified location first. If they reach the specified location within 5 minutes (i.e., before the event ends) they receive 1 of each potion.

This example event is a guise for an egress simulation. The event map is based off the floor plan of a building, and the goal locations are exits. As long as the rewards system is carefully considered, any number of other factors may be added in. For example, if we wish to determine the effect of handicapped individuals on egress, we might arbitrarily select users and slow their maximum speeds. Correspondingly, we should reward them with *two* of each potion if they complete the event.

Validation of unannounced MMOHILS will be discussed in Section 3.3, along with that for experimental MMOHILS. However, it is important to note that data from unannounced MMOHILS will inherently contain much more noise than that from the experimental version. Those who merely wish to simulate massive groups of users without the need for an *un-announced* aspect might consider simply gathering experiment participants *en masse* and running a large-scale experimental MMOHILS.

3.3 Validating a MMOHILS

Validation techniques for MMOHILS borrow heavily from traditional simulation validation and human-in-the-loop simulation validation. In particular, the **physical environ-**

ment, the agent's **capabilities** within this environment (over time), the **psychological state** of each human user, and the **role** that user plays in the environment must all be validated to the necessary degree of precision demanded by the simulation's developers. We will cover the relevant techniques out-of-order.

Traditional black box validation of a pedestrian agent-based simulation is concerned with comparing the fluctuations in agent densities over the course of the simulation with those of actual observed humans. (Berrou 2000) Moreover, the size and shape of all bottlenecks are validated with respect to observed bottlenecks. These validation steps must also be performed for a MMOHILS with pedestrian elements.

Each agent's capabilities must also remain valid as the simulation progresses. Smoke, for example, obscures vision and reduces an agent's average speed. A dosage metric (such as Fractional Effective Dose (Hartzell 1988)) can be used to adjust users' capabilities in the simulated world; moreover, where more than one metric exist, one might be used to adjust while the other is used to validate these adjustments as the simulation progresses. Static properties are fairly simple to validate; Berrou, for example, determines the distributions of physical *radius*, *preferred free speed*, and *personal space* which pedestrians in the Legion model conform to. Similar distributions exist (or can be gleaned from observation) for most domains of interest.

In a given simulation run, relevant psychological factors must also be validated on a per-user basis. For example, if the simulation intends to test evacuation under stressful conditions, then users must be surveyed *post factum* so as to determine their level of perceived stress. For unannounced MMOHILS, this may seem difficult. Internet game players are notoriously obstreperous and prone to capricious or oppositional-defiant assertion. Currently, various psychological tests exist which hide the objective of the test from the agent being questioned. Seemingly-innocuous questions can be used to gain insight into the user's internal state. Consider the "Object Identification Test", which presents test-takers with a series of ambiguous objects that steadily come to resemble everyday items. (Harris 1967) Users under stress take longer to recognize these objects; we might fit such a test into our example game world by telling users to "identify this object" to claim their reward.

Validating each agent's role in the world is the biggest challenge MMOHILS faces. While rule bases can be mined automatically from very large data sets, there's no guarantee that these rules are nothing more than coincidence. Expert systems, on the other hand, are far too expensive. Hopkinson and Sepulveda (1995) proposed using case-based reasoning to validate the simulation in real-time. A small number of *representative* cases are manually entered into the case base. These cases describe invalid user behavior, and any user who acts similarly is warned (and his data marked invalid). Borderline cases are tagged for (human) review later, and may then be entered into the case base. A threshold can be set for what constitutes "similar" behavior and "borderline" cases. Setting this threshold relatively high will catch more Type-II errors, but commit more Type-I errors. The nature of the simulation will determine which error type is least damaging to one's data.

Finally, the environment in which the test takes place must itself be validated as a proper representation of the actual environment. For example, a simulation of a burning building would not be valid if floors and walls did not collapse as the fire progressed. For this, Turing-test validation is probably best: an expert should be brought in and shown representations of users moving through the simulation and, comparatively, users moving through a real-world example, and he should be incapable of distinguishing the simulation from the real data.

"The modeler is most concerned with how accurately and with what degree of precision… the developed computer simulation model represent[s] the real world." (Hopkinson and Sepulveda 1995) In certain cases, the desired level of precision may be impossible to achieve with MMOHILS. For example, determining that your users are behaving psychologically identical to humans in a burning building may require monumental advances in brain imaging technology. However, determining that said users are as stressed and disoriented as their real-world counterparts is much easier, and may provide sufficient precision to validate the model and each simulation run.

4 DISCUSSION

4.1 Benefits of MMOHILS

This unique approach to simulation has many benefits:

- The human components of a simulation are usually the most difficult to validate; if well-researched AI does not exist in a domain, MMOHILS can be deployed with low cost and high accuracy.
- In domains well understood from an AI perspective, MMOHILS provides a means of validating AI-based agents or of training them for new situations and conditions.
- Although our approach is to design a MMOHILS from the ground up, there is no argument against adapting an existing MMO game (such as Second Life) to serve this purpose. In this case, great care would have to be taken to validate all algorithms which distort the model (e.g., time-warp).
- Valid AI agents may exhibit emergent behavior which is unacceptable; MMOHILS scales well with respect to the number of agents.

- Although human involvement may increase the cost (in time and money) of simulation, the real-time verification technique introduced allows the simulation to detect small errors in users' decisions before they become significant. In some cases, steps may be taken to inform the user of his mistake, and thereby salvage his data in whole or in part.

4.2 Roadmap

We are currently in the early stages of implementing an experimental MMOHILS as a proof-of-concept. Following this, an unannounced MMOHILS will be designed (and likely deployed) and a framework for simulations of this kind will be formally defined.

5 CONCLUSION

Massively Multiplayer Online Human-In-the-Loop Simulation was presented as a venue of future research in simulation and gaming. Two systems were clearly laid out: one for experimental procedures, and one for un-announced observational studies. Both can be classified as discrete time-stepped simulations. Standard validation techniques can be applied to MMOHILS; in addition, a real-time technique for validating users' roles was referenced.

MMOHILS represents a simple solution to the difficult problem of modeling human behavior. It increases the accuracy of simulations of this nature, and broadens the range of applications of experimental research.

ACKNOWLEDGMENTS

This work was supported by the MOE AcRF Tier 1 grant R-252-000-314-112.

REFERENCES

Berrou, J., J. Beecham, P. Quaglia, M. Kagarlis, and A. Gerodimos. *Calibration and Validation of the Legion Simulation Model Using Empirical Data*. 2000. PhD Thesis, University of Warwick.

Chaturvedi, A., A. Mellema, S. Filatyev, and J. Gore. 2006. *DDDAS for Fire and Agent Evacuation Modeling of the Rhode Island Nightclub Fire*. ICCS2006 Workshop on Dynamic Data Driven Applications Systems.

Daamen, W., and S. P. Hoogendoorn. 2003. *Qualitative results from pedestrian laboratory experiments.* In E.R. Galea (Ed.), Pedestrian and evacuation dynamics 2003 (pp. 121-132). London: CMS Press.

Fruin, J. 1987. *Pedestrian Planning and Design, 2nd Edition*. Elevator World, Mobile AL.

Harris, W. 1967. *The Object Identification Test: A Stress-Sensitive Perceptual Test.* Human Factors Research Inc. Goleta, California.

Hartzell, G. 1998. *The Fractional Effective Dose Model for Assessment of Toxic Hazards in Fires*. Journal of Fire Sciences, 6(5): 356-362.

Hopkinson, W., and J. Sepulveda. 1995. Real Time Validation of Man-in-the-Loop Simulations. In *Proceedings of the 1995 Winter Simulation Conference*, eds. C. Alexopoulos, and K. Kang, 1250-1256. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Jain, S., and C. McLean. 2006. A Concept Prototype for Integrated Gaming and Simulation for Incident Management. In *Proceedings of the 2006 Winter Simulation Conference*. eds. L. Felipe Perrone, B. Lawson, J. Liu, and F. P. Wieland, 493-500. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Okazaki, S. and M. Satoshi. 1993. *A Study of Simulation Model for Pedestrian Movement with Evacuation and Queuing*. In: Proc. International Conference on Engineering for Crowd Safety.

Pan, X., C. Han, and K. Law. *A Multi-Agent Based Simulation Framework for the Simulation of Human and Social Behaviors in Egress Analysis*. The International Conference on Computing in Civil Engineering.

Santos, G., and B. Aguirre. 2004. *A Critical Review of Emergency Evacuation Simulation Models*. NIST Workshop on Building Occupant Movement during Fire Emergencies.

Tay, V. 2005. *Massively Multiplayer Online Game (MMOG) — a Proposed Approach for Military Applications*. International Conference on Cyberworlds, IEEE, 396-400.

Tierney, B., W. Johnston, H. Herzog, G. Hoo, G. Jin, J. Lee, L. Chen, and D. Rotem. 1994. *Using High Speed Networks to Enable Distributed Parallel Image Server Systems*. IEEE.

AUTHOR BIOGRAPHIES

SETH HETU is working towards his PhD in Computer Science at the National University of Singapore. His focus is in Crisis Management Simulation, and in Symbiotic Simulation. E- mail: `<seth.hetu@gmail.com>`

GARY TAN is an Associate Professor at the Department of Computer Science, School of Computing, National University of Singapore. His current research interests are Parallel and Distributed Simulation, Model Composability, Crisis Management Simulation and Symbiotic Simulation. E- mail: `<gtan@comp.nus.edu.sg>`